Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	62501	table same (update or updated or updating)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 20:55
L2	384400	CD or CDRW or CD-RW or dvd or dvd+RW or dvd-RW	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 20:49
L3	719241	manag$5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 20:49
L4	50055	L3 with table	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 20:49
L5	15826	L1 and L4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 20:49
L6	4462	L2 and L5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 20:49
L7	674	L4 with L2	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 21:07
L8	139	L1 and L7	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 20:49
L9	862	711/156.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 20:49
L10	779	711/111.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 20:49
L11	184	711/221.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 20:49

L12	1798	L9 or L10 or L11	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 20:49
L16	21934749	@ad<"20020210"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 20:51
L17	1	8 and 12	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 20:52
L18	1	16 and 17	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 20:52
L19	1	("6389509").URPN.	USPAT	OR	ON	2006/01/17 21:02
L20	7	table with (update or updated or updating) with (use or used or using or usage) with L2	USPAT	OR	ON	2006/01/17 20:55
L21	6	16 and 20	USPAT	OR	ON	2006/01/17 20:55
L22	6	16 and 20	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 20:55
L23	14	table with (update or updated or updating) with (use or used or using or usage) with L2	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 20:56
L24	10	16 and 23	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 21:00
L25	2415	usage near5 status	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 21:04
L26	22	2 same 25	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 21:00
L27	14	1 and 26	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 21:00

L28	4	16 and 27	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 21:00
L29	590	(reconstruct or reconstructed or reconstructing) near5 table	USPAT	OR	ON	2006/01/17 21:04
L30	5	25 and 29	USPAT	OR	ON	2006/01/17 21:03
L31	5	16 and 30	USPAT	OR	ON	2006/01/17 21:04
L32	5	16 and 30	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 21:04
L33	1100	(reconstruct or reconstructed or reconstructing) near5 table	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 21:04
L34	8	25 and 33	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 21:04
L35	3	34 not 32	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 21:05
L36	0	16 and 35	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 21:06
L37	283	711/156.ccls.	US-PGPUB	OR	ON	2006/01/17 21:06
L38	56	16 and 37	US-PGPUB	OR	ON	2006/01/17 21:06
L39	1	33 and 38	US-PGPUB	OR	ON	2006/01/17 21:06
L40	0	6 and 38	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 21:07
L41	0	7 and 38	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/17 21:07

Results for "((disk <and> table <and> management)<in>metadata)"

e-mail printer friendly

Your search matched 13 of 1302021 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.

» Search Options

View Session History

New Search

Modify Search

((disk <and> table <and> management)<in>metadata)

Check to search only within this results set

Display Format: Citation Citation & Abstract

» Key

IEEE JNL	IEEE Journal or Magazine
IEE JNL	IEE Journal or Magazine
IEEE CNF	IEEE Conference Proceeding
IEE CNF	IEE Conference Proceeding
IEEE STD	IEEE Standard

Select Article Information

1. **Disk read-write optimizations and data integrity in transaction systems using write-ahead logging**
Mohan, C.;
Data Engineering, 1995. Proceedings of the Eleventh International Conference on
6-10 March 1995 Page(s):324 - 331
Digital Object Identifier 10.1109/ICDE.1995.380378
AbstractPlus | Full Text: PDF(840 KB) IEEE CNF

2. **A multi-resolution block storage model for database design**
Jingren Zhou; Ross, K.A.;
Database Engineering and Applications Symposium, 2003. Proceedings. Seventh International
16-18 July 2003 Page(s):22 - 31
AbstractPlus | Full Text: PDF(352 KB) IEEE CNF

3. **Dynamic power management for nonstationary service requests**
Eui-Young Chung; Benini, L.; Bogliolo, A.; Yung-Hsiang Lu; De Micheli, G.;
Computers, IEEE Transactions on
Volume 51, Issue 11, Nov. 2002 Page(s):1345 - 1361
Digital Object Identifier 10.1109/TC.2002.1047758
AbstractPlus | References | Full Text: PDF(804 KB) IEEE JNL

4. **Efficient integration of data mining techniques in database management systems**
Bentayeb, F.; Darmont, J.; Udrea, C.;
Database Engineering and Applications Symposium, 2004. IDEAS '04. Proceedings. International
7-9 July 2004 Page(s):59 - 67
Digital Object Identifier 10.1109/IDEAS.2004.1319778
AbstractPlus | Full Text: PDF(325 KB) IEEE CNF

5. **Operating system management of large partially allocated objects**
Hollins, M.; Rosenberg, J.;
Object Orientation in Operating Systems, 1993., Proceedings of the Third International Workshop on
9-10 Dec. 1993 Page(s):42 - 50
Digital Object Identifier 10.1109/IWOOOS.1993.324928
AbstractPlus | Full Text: PDF(856 KB) IEEE CNF

6. **Persistent caching: an implementation technique for complex objects with object identity**
Kato, K.; Masuda, T.;
Software Engineering, IEEE Transactions on
Volume 18, Issue 7, July 1992 Page(s):631 - 645
Digital Object Identifier 10.1109/32.148481
AbstractPlus | Full Text: PDF(1376 KB) IEEE JNL

7. **A multi-stage retrieval algorithm for postings on index tables**
Fuketa, M.; Tsuji, T.; Ando, K.; Aoe, J.;
Systems, Man, and Cybernetics, 1998. 1998 IEEE International Conference on
Volume 5, 11-14 Oct. 1998 Page(s):4578 - 4583 vol.5
Digital Object Identifier 10.1109/ICSMC.1998.727573

AbstractPlus | Full Text: PDF(544 KB) IEEE CNF

8. **Performance gains using subpages for cache coherency control**
Listl, A.; Bozas, G.;
Database and Expert Systems Applications, 1997. Proceedings., Eighth International Workshop on
1-2 Sept. 1997 Page(s):754 - 760
Digital Object Identifier 10.1109/DEXA.1997.617421

AbstractPlus | Full Text: PDF(612 KB) IEEE CNF

9. **Quality-controlled compression of sets of images**
Livny, M.; Ratnakar, V.;
Multimedia Database Management Systems, 1996., Proceedings of International Workshop on
14-16 Aug. 1996 Page(s):109 - 114
Digital Object Identifier 10.1109/MMDBMS.1996.541861

AbstractPlus | Full Text: PDF(800 KB) IEEE CNF

10. **MIDAS: design philosophy and internals**
Pyung-Chul Kim; Hwan Ik Choi; Yoon-Joon Lee; Sang Ho Lee; Myung-Joon Kim;
Computers and Communications, 1992. Conference Proceedings., Eleventh Annual International Phoenix Conference on
1-3 April 1992 Page(s):132 - 139
Digital Object Identifier 10.1109/PCCC.1992.200549

AbstractPlus | Full Text: PDF(728 KB) IEEE CNF

11. **An efficient multiprocessor merge algorithm**
Varman, P.J.; Iyer, B.R.; Haderle, D.J.;
Databases, Parallel Architectures and Their Applications,. PARBASE-90, International Conference on
7-9 March 1990 Page(s):276 - 283
Digital Object Identifier 10.1109/PARBSE.1990.77152

AbstractPlus | Full Text: PDF(600 KB) IEEE CNF

12. **Design and implementation of the concurrency control manager in the main-memory DBMS Tachyon**
Sang-Wook Kim; Wan Choi; Byoung-Ho Kim;
Computer Software and Applications Conference, 2002. COMPSAC 2002. Proceedings. 26th Annual International
26-29 Aug. 2002 Page(s):635 - 641
Digital Object Identifier 10.1109/CMPSAC.2002.1045075

AbstractPlus | Full Text: PDF(1622 KB) IEEE CNF

13. **Trends in automating physical database design**
Zilio, D.C.; Lightstone, S.; Lohman, G.M.;
Industrial Informatics, 2003. INDIN 2003. Proceedings. IEEE International Conference on
21-24 Aug. 2003 Page(s):441 - 445
Digital Object Identifier 10.1109/INDIN.2003.1300376

AbstractPlus | Full Text: PDF(1599 KB) IEEE CNF

Indexed by Inspec

© Copyright 2005 IEEE – All Rights Reserved

Results for "((((disk <and> table <and> management)<in>metadata))<and>(use <or> usa..." e-mail printer friendly

Your search matched **10** of **13** documents.

A maximum of **100** results are displayed, **25** to a page, sorted by **Relevance** in **Descending** order.

» Search Options

View Session History

New Search

Modify Search

((((disk <and> table <and> management)<in>metadata))<and>(use <or> usage<in>m >>

Check to search only within this results set

Display Format: Citation Citation & Abstract

» Key

IEEE JNL	IEEE Journal or Magazine
IEE JNL	IEE Journal or Magazine
IEEE CNF	IEEE Conference Proceeding
IEE CNF	IEE Conference Proceeding
IEEE STD	IEEE Standard

Select Article Information

1. **A multi-resolution block storage model for database design**
Jingren Zhou; Ross, K.A.;
Database Engineering and Applications Symposium, 2003. Proceedings. Seventh International
16-18 July 2003 Page(s):22 - 31
AbstractPlus | Full Text: PDF(352 KB) IEEE CNF

2. **Dynamic power management for nonstationary service requests**
Eui-Young Chung; Benini, L.; Bogliolo, A.; Yung-Hsiang Lu; De Micheli, G.;
Computers, IEEE Transactions on
Volume 51, Issue 11, Nov. 2002 Page(s):1345 - 1361
Digital Object Identifier 10.1109/TC.2002.1047758
AbstractPlus | References | Full Text: PDF(804 KB) IEEE JNL

3. **Efficient integration of data mining techniques in database management systems**
Bentayeb, F.; Darmont, J.; Udrea, C.;
Database Engineering and Applications Symposium, 2004. IDEAS '04. Proceedings. International
7-9 July 2004 Page(s):59 - 67
Digital Object Identifier 10.1109/IDEAS.2004.1319778
AbstractPlus | Full Text: PDF(325 KB) IEEE CNF

4. **A multi-stage retrieval algorithm for postings on index tables**
Fuketa, M.; Tsuji, T.; Ando, K.; Aoe, J.;
Systems, Man, and Cybernetics, 1998. 1998 IEEE International Conference on
Volume 5, 11-14 Oct. 1998 Page(s):4578 - 4583 vol.5
Digital Object Identifier 10.1109/ICSMC.1998.727573
AbstractPlus | Full Text: PDF(544 KB) IEEE CNF

5. **Disk read-write optimizations and data integrity in transaction systems using write-ahead logging**
Mohan, C.;
Data Engineering, 1995. Proceedings of the Eleventh International Conference on
6-10 March 1995 Page(s):324 - 331
Digital Object Identifier 10.1109/ICDE.1995.380378
AbstractPlus | Full Text: PDF(840 KB) IEEE CNF

6. **An efficient multiprocessor merge algorithm**
Varman, P.J.; Iyer, B.R.; Haderle, D.J.;
Databases, Parallel Architectures and Their Applications,. PARBASE-90, International Conference on
7-9 March 1990 Page(s):276 - 283
Digital Object Identifier 10.1109/PARBSE.1990.77152
AbstractPlus | Full Text: PDF(600 KB) IEEE CNF

7. **Performance gains using subpages for cache coherency control**
Listl, A.; Bozas, G.;
Database and Expert Systems Applications, 1997. Proceedings., Eighth International Workshop on
1-2 Sept. 1997 Page(s):754 - 760
Digital Object Identifier 10.1109/DEXA.1997.617421

AbstractPlus | Full Text: PDF(612 KB) IEEE CNF

8. **Trends in automating physical database design**
Zilio, D.C.; Lightstone, S.; Lohman, G.M.;
Industrial Informatics, 2003. INDIN 2003. Proceedings. IEEE International Conference on
21-24 Aug. 2003 Page(s):441 - 445
Digital Object Identifier 10.1109/INDIN.2003.1300376

AbstractPlus | Full Text: PDF(1599 KB) IEEE CNF

9. **Quality-controlled compression of sets of images**
Livny, M.; Ratnakar, V.;
Multimedia Database Management Systems, 1996., Proceedings of International Workshop on
14-16 Aug. 1996 Page(s):109 - 114
Digital Object Identifier 10.1109/MMDBMS.1996.541861

AbstractPlus | Full Text: PDF(800 KB) IEEE CNF

10. **Operating system management of large partially allocated objects**
Hollins, M.; Rosenberg, J.;
Object Orientation in Operating Systems, 1993., Proceedings of the Third International Workshop on
9-10 Dec. 1993 Page(s):42 - 50
Digital Object Identifier 10.1109/IWOOOS.1993.324928

AbstractPlus | Full Text: PDF(856 KB) IEEE CNF

Help Contact Us Privacy & Security IEEE.org

Indexed by Inspec

© Copyright 2005 IEEE – All Rights Reserved

Results for "((((((disk <and> table <and> management)<in>metadata))<and>(use <or> u..."

e-mail printer friendly

Your search matched 0 of 10 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by **Relevance** in **Descending** order.

» Search Options

View Session History

New Search

Modify Search

((((((disk <and> table <and> management)<in>metadata))<and>(use <or> usage<in>

Check to search only within this results set

Display Format: Citation Citation & Abstract

» Key

IEEE JNL	IEEE Journal or Magazine
IEE JNL	IEE Journal or Magazine
IEEE CNF	IEEE Conference Proceeding
IEE CNF	IEE Conference Proceeding
IEEE STD	IEEE Standard

No results were found.

Please edit your search criteria and try again. Refer to the Help pages if you need assistance revising your search.

Indexed by Inspec

© Copyright 2005 IEEE – All Rights Reserved

Edit an existing query or compose a new query in the Search Query Display.

Select a search number (#) to:

- Add a query to the Search Query Display
- Combine search queries using AND, OR, or NOT
- Delete a search
- Run a search

Tue, 17 Jan 2006, 9:17:15 PM EST

Search Query Display

[redacted]

[redacted]

Recent Search Queries		Results
#1	((disk <and> table <and> management)<in>metadata)	13
#2	((((disk <and> table <and> management)<in>metadata))<AND>(use <or> usage<in>metadata))	10
#3	((((((disk <and> table <and> management)<in>metadata))<and>(use <or> usage<in>metadata)))<AND>(~~disk management~~<in>metadata))	0

[redacted]

Indexed by Inspec

© Copyright 2005 IEEE – All Rights Reserved

Terms used **disk management** and **table** and **usage** Found **16,183** of **169,166**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

☐ Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown Relevance scale

1 A study of I/O behavior of perfect benchmarks on a multiprocessor

A. L. Narasimha Reddy, Prithviraj Banerjee

May 1990 **ACM SIGARCH Computer Architecture News , Proceedings of the 17th annual international symposium on Computer Architecture ISCA '90**, Volume 18 Issue 3a

Publisher: ACM Press

Full text available: pdf(984.37 KB) Additional Information: full citation, abstract, references, citings, index terms

The I/O behavior of some scientific applications, a subset of Perfect benchmarks, executing on a multiprocessor is studied. The aim of this study is to explore the various patterns of I/O access of large scientific applications and to understand the impact of this observed behavior on the I/O subsystem architecture. I/O behavior of the program is characterized by the demands it imposes on the I/O subsystem. It is observed that implicit I/O or paging is not a major problem for the applicatio ...

2 The logical disk: a new approach to improving file systems

Wiebren de Jonge, M. Frans Kaashoek, Wilson C. Hsieh

December 1993 **ACM SIGOPS Operating Systems Review , Proceedings of the fourteenth ACM symposium on Operating systems principles SOSP '93**, Volume 27 Issue 5

Publisher: ACM Press

Full text available: pdf(1.55 MB) Additional Information: full citation, abstract, references, citings, index terms

The Logical Disk (LD) defines a new interface to disk storage that separates file management and disk management by using logical block numbers and block lists. The LD interface is designed to support multiple file systems and to allow multiple implementations, both of which are important given the increasing use of kernels that support multiple operating system personalities.A log-structured implementation of LD (LLD) demonstrates that LD can be implemented efficiently. LLD adds about 5% to 10% ...

Keywords: MINIX, UNIX, disk storage management, file system organization, file system performance, high write performance, log-structured file system, logical disk

3 Partially preemptible hash joins

Hwee Hwa Pang, Michael J. Carey, Miron Livny

June 1993 **ACM SIGMOD Record , Proceedings of the 1993 ACM SIGMOD international conference on Management of data SIGMOD '93**, Volume 22 Issue 2

Publisher: ACM Press

Full text available: pdf(1.42 MB) Additional Information: full citation, abstract, references, citings, index terms

With the advent of real-time and goal-oriented database systems, priority scheduling is likely to be an important feature in future database management systems. A consequence of priority scheduling is that a transaction may lose its buffers to higher-priority transactions, and may be given additional memory when transactions leave the system. Due to their heavy reliance on main memory, hash joins are especially vulnerable to fluctuations in memory availability. Previous studies have propose ...

4 The simulation of time sharing systems

Norman R. Nielsen

July 1967 **Communications of the ACM**, Volume 10 Issue 7

Publisher: ACM Press

Full text available: pdf(2.23 MB) Additional Information: full citation, abstract, references, citings, index terms

The development of new large scale time-sharing systems has raised a number of problems for computation center management. Not only is it necessary to develop an appropriate hardware configuration for these systems, but appropriate software adjustments must be made. Unfortunately, these systems often do not respond to changes in the manner that intuition would suggest, and there are few guides to assist in the analysis of performance characteristics. The development of a comprehensive simul ...

5 Energy-aware computing: Exposing disk layout to compiler for reducing energy consumption of parallel disk based systems

S. W. Son, G. Chen, M. Kandemir, A. Choudhary

June 2005 **Proceedings of the tenth ACM SIGPLAN symposium on Principles and practice of parallel programming**

Publisher: ACM Press

Full text available: pdf(668.20 KB) Additional Information: full citation, abstract, references, index terms

Disk subsystem is known to be a major contributor to overall power consumption of high-end parallel systems. Past research proposed several architectural level techniques to reduce disk power by taking advantage of idle periods experienced by disks. While such techniques have been known to be effective in certain cases, they share a common drawback: they operate in a reactive manner; i.e., they control disk power by observing past disk activity (e.g., idle and active periods) and estimating futu ...

Keywords: low power, optimizing compiler, parallel disk

6 Parameterized Ray-tracing

C. H. Séquin, E. K. Smyrl

July 1989 **ACM SIGGRAPH Computer Graphics , Proceedings of the 16th annual conference on Computer graphics and interactive techniques SIGGRAPH '89**, Volume 23 Issue 3

Publisher: ACM Press

Full text available: pdf(4.62 MB) Additional Information: full citation, abstract, references, citings, index terms

The construction and refinement of a computer graphics scene is unacceptably slow when using ray tracing. We introduce a new technique to speed up the generation of successive ray traced images when the geometry of the scene remains constant and only the light source intensities and the surface properties need to be adjusted. When the scene is first ray traced, an expression parameterized in the color of all lights and the surface property coefficients of all objects is calculated and stored for ...

7 A modular simulation of TSS/360

John W. McCredie, Steven J. Schlesinger

December 1970 **Proceedings of the fourth annual conference on Applications of simulation**

Publisher: Winter Simulation Conference

Full text available: pdf(480.18 KB) Additional Information: full citation, abstract, references, citings

This paper describes a modular simulation of TSS/360, a virtual memory time-sharing computer system. A detailed model of the table driven scheduler is embedded in an environment consisting of higher level models of other subsystems. The original goal of the project was to provide a useful tool to aid in the tuning of the monitor's scheduling algorithm. As the simulation evolved, the model served as a vehicle to study proposed hardware changes.

8 RAID: high-performance, reliable secondary storage

Peter M. Chen, Edward K. Lee, Garth A. Gibson, Randy H. Katz, David A. Patterson

June 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 2

Publisher: ACM Press

Full text available: pdf(3.60 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Disk arrays were proposed in the 1980s as a way to use parallelism between multiple disks to improve aggregate I/O performance. Today they appear in the product lines of most major computer manufacturers. This article gives a comprehensive overview of disk arrays and provides a framework in which to organize current and future work. First, the article introduces disk technology and reviews the driving forces that have popularized disk arrays: performance and reliability. It discusses the tw ...

Keywords: RAID, disk array, parallel I/O, redundancy, storage, striping

9 A study of initialization in Linux and OpenBSD

Catherine Dodge, Cynthia Irvine, Thuy Nguyen

April 2005 **ACM SIGOPS Operating Systems Review**, Volume 39 Issue 2

Publisher: ACM Press

Full text available: pdf(2.02 MB) Additional Information: full citation, abstract, references, index terms

The code that initializes a system can be notoriously difficult to understand. In secure systems, initialization is critical for establishing a starting state that is secure. This paper explores two architectures used for bringing an operating system to its initial state, once the operating system gains control from the boot loader. Specifically, the ways in which the OpenBSD and Linux operating systems handle initialization are dissected.

10 Computing curricula 2001

September 2001 **Journal on Educational Resources in Computing (JERIC)**

Publisher: ACM Press

Full text available: pdf(613.63 KB) html(2.78 KB) Additional Information: full citation, references, citings, index terms

11 Caching in the Sprite network file system

Michael N. Nelson, Brent B. Welch, John K. Ousterhout

February 1988 **ACM Transactions on Computer Systems (TOCS)**, Volume 6 Issue 1

Publisher: ACM Press

Full text available: pdf(1.68 MB) Additional Information: full citation, abstract, references, citings, index terms, review

The Sprite network operating system uses large main-memory disk block caches to achieve high performance in its file system. It provides non-write-through file caching on both client and server machines. A simple cache consistency mechanism permits files to be shared by multiple clients without danger of stale data. In order to allow the file cache to

occupy as much memory as possible, the file system of each machine negotiates with the virtual memory system over physical memory usage and c ...

12 Query evaluation techniques for large databases

Goetz Graefe
June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2
Publisher: ACM Press

Full text available: pdf(9.37 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

13 A structural view of the Cedar programming environment

Daniel C. Swinehart, Polle T. Zellweger, Richard J. Beach, Robert B. Hagmann
August 1986 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 8 Issue 4
Publisher: ACM Press

Full text available: pdf(6.32 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper presents an overview of the Cedar programming environment, focusing on its overall structure—that is, the major components of Cedar and the way they are organized. Cedar supports the development of programs written in a single programming language, also called Cedar. Its primary purpose is to increase the productivity of programmers whose activities include experimental programming and the development of prototype software systems for a high-performance personal computer. T ...

14 A resource management framework for priority-based physical-memory allocation

Kingsley Cheung, Gernot Heiser
January 2002 **Australian Computer Science Communications , Proceedings of the seventh Asia-Pacific conference on Computer systems architecture - Volume 6 CRPITS '02**, Volume 24 Issue 3
Publisher: Australian Computer Society, Inc. , IEEE Computer Society Press
Full text available: pdf(1.32 MB) Additional Information: full citation, abstract, references, index terms

Most multitasking operating systems support scheduling priorities in order to ensure that processor time is allocated to important or time-critical processes in preference to less important ones. Ideally this would prevent a low-priority process from slowing the execution of a high-priority one. In practice, strict prioritisation is undermined by a lack of suitable allocation policy for resources other than CPU time. For example, a low priority process may degrade the execution speed of a high-p ...

15 Techniques for efficient inline tracing on a shared-memory multiprocessor

S. J. Eggers, David R. Keppel, Eric J. Koldinger, Henry M. Levy
April 1990 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1990 ACM SIGMETRICS conference on Measurement and modeling of computer systems SIGMETRICS '90**, Volume 18 Issue 1
Publisher: ACM Press

Full text available: Additional Information: full citation, abstract, references, citings, index

pdf(1.12 MB) terms

While much current research concerns multiprocessor design, few traces of parallel programs are available for analyzing the effect of design trade-offs. Existing trace collection methods have serious drawbacks: trap-driven methods often slow down program execution by more than 1000 times, significantly perturbing program behavior; microcode modification is faster, but the technique is neither general nor portable. This paper describes a new tool, called MPTRACE, for collecting tr ...

16 Parity declustering for continuous operation in redundant disk arrays

Mark Holland, Garth A. Gibson

September 1992 **ACM SIGPLAN Notices , Proceedings of the fifth international conference on Architectural support for programming languages and operating systems ASPLOS-V**, Volume 27 Issue 9

Publisher: ACM Press

Full text available: pdf(1.57 MB) Additional Information: full citation, references, citings, index terms

17 Application performance and flexibility on exokernel systems

M. Frans Kaashoek, Dawson R. Engler, Gregory R. Ganger, Hector M. Briceño, Russell Hunt, David Mazières, Thomas Pinckney, Robert Grimm, John Jannotti, Kenneth Mackenzie

October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles SOSP '97**, Volume 31 Issue 5

Publisher: ACM Press

Full text available: pdf(2.39 MB) Additional Information: full citation, references, citings, index terms

18 ECOSystem: managing energy as a first class operating system resource

Heng Zeng, Carla S. Ellis, Alvin R. Lebeck, Amin Vahdat

October 2002 **ACM SIGPLAN Notices , ACM SIGARCH Computer Architecture News , ACM SIGOPS Operating Systems Review , Proceedings of the 10th international conference on Architectural support for programming languages and operating systems ASPLOS-X**, Volume 37 , 30 , 36 Issue 10 , 5 , 5

Publisher: ACM Press

Full text available: pdf(1.17 MB) Additional Information: full citation, abstract, references, citings

Energy consumption has recently been widely recognized as a major challenge of computer systems design. This paper explores how to support energy as a first-class operating system resource. Energy, because of its global system nature, presents challenges beyond those of conventional resource management. To meet these challenges we propose the Currentcy Model that unifies energy accounting over diverse hardware components and enables fair allocation of available energy among applications. Our par ...

19 Extended ephemeral logging: log storage management for applications with long lived transactions

John S. Keen, William J. Dally

March 1997 **ACM Transactions on Database Systems (TODS)**, Volume 22 Issue 1

Publisher: ACM Press

Full text available: pdf(566.34 KB) Additional Information: full citation, references, index terms, review

Keywords: OLTP, disk management, logging, long transactions

20 Safe and efficient sharing of persistent objects in Thor

B. Liskov, A. Adya, M. Castro, S. Ghemawat, R. Gruber, U. Maheshwari, A. C. Myers, M. Day,

L. Shrira
June 1996 **ACM SIGMOD Record , Proceedings of the 1996 ACM SIGMOD international conference on Management of data SIGMOD '96**, Volume 25 Issue 2
Publisher: ACM Press

Full text available: pdf(1.42 MB) Additional Information: full citation, abstract, references, citings, index terms

Thor is an object-oriented database system designed for use in a heterogeneous distributed environment. It provides highly-reliable and highly-available persistent storage for objects, and supports safe sharing of these objects by applications written in different programming languages.Safe heterogeneous sharing of long-lived objects requires encapsulation: the system must guarantee that applications interact with objects only by invoking methods. Although safety concerns are important, most obj ...

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.
Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player

Terms used **disk management** and **table** and **disk usage** Found **596** of **169,166**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

☐ Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown Relevance scale

1 Energy-aware computing: Exposing disk layout to compiler for reducing energy consumption of parallel disk based systems

S. W. Son, G. Chen, M. Kandemir, A. Choudhary

June 2005 **Proceedings of the tenth ACM SIGPLAN symposium on Principles and practice of parallel programming**

Publisher: ACM Press

Full text available: pdf(668.20 KB) Additional Information: full citation, abstract, references, index terms

Disk subsystem is known to be a major contributor to overall power consumption of high-end parallel systems. Past research proposed several architectural level techniques to reduce disk power by taking advantage of idle periods experienced by disks. While such techniques have been known to be effective in certain cases, they share a common drawback: they operate in a reactive manner; i.e., they control disk power by observing past disk activity (e.g., idle and active periods) and estimating futu ...

Keywords: low power, optimizing compiler, parallel disk

2 A survey of power management techniques in mobile computing operating systems

Gregory F. Welch

October 1995 **ACM SIGOPS Operating Systems Review**, Volume 29 Issue 4

Publisher: ACM Press

Full text available: pdf(763.75 KB) Additional Information: full citation, abstract, citings, index terms

Many factors have contributed to the birth and continued growth of mobile computing, including recent advances in hardware and communications technology. With this new paradigm however come new challenges in computer operating systems development. These challenges include heretofore relatively unusual items such as frequent network disconnections, communications bandwidth limitations, resource restrictions, and power limitations. It is the last of these challenges that we shall explore in this p ...

3 The logical disk: a new approach to improving file systems

Wiebren de Jonge, M. Frans Kaashoek, Wilson C. Hsieh

December 1993 **ACM SIGOPS Operating Systems Review , Proceedings of the fourteenth ACM symposium on Operating systems principles SOSP '93**, Volume 27 Issue 5

Publisher: ACM Press

Full text available: pdf(1.55 MB) Additional Information: full citation, abstract, references, citings, index terms

The Logical Disk (LD) defines a new interface to disk storage that separates file management and disk management by using logical block numbers and block lists. The LD

interface is designed to support multiple file systems and to allow multiple implementations, both of which are important given the increasing use of kernels that support multiple operating system personalities.A log-structured implementation of LD (LLD) demonstrates that LD can be implemented efficiently. LLD adds about 5% to 10% ...

Keywords: MINIX, UNIX, disk storage management, file system organization, file system performance, high write performance, log-structured file system, logical disk

4 Linux Gazette: Disk Hog: Tracking System Disk Usage

Ivan Griffin
December 1997 **Linux Journal**
Publisher: Specialized Systems Consultants, Inc.
Full text available: html(8.33 KB) Additional Information: full citation, index terms

5 Session 3: Scalability and resource usage of an OLAP benchmark on clusters of PCs

Michela Taufer, Thomas Stricker, Roger Weber
August 2002 **Proceedings of the fourteenth annual ACM symposium on Parallel algorithms and architectures**
Publisher: ACM Press
Full text available: pdf(219.90 KB) Additional Information: full citation, abstract, references, index terms

Designing clusters of PCs for distributed databases processing OLAP(On Line Analytical Processing) workloads in parallel with good scalability remains a particular challenge as we are lacking a deep understanding of the architectural issues around resource usage by standard DBMSs on distributed platforms.To address this problem, we present a novel performance monitoring framework for filtering and abstracting samples of performance data from low level counters into a high level performance pictu ...

Keywords: cluster of PCs, distributed OLAP processing, parallel databases, performance analysis, workload characterization

6 Continuous profiling: where have all the cycles gone?

Jennifer M. Anderson, Lance M. Berc, Jeffrey Dean, Sanjay Ghemawat, Monika R. Henzinger, Shun-Tak A. Leung, Richard L. Sites, Mark T. Vandevoorde, Carl A. Waldspurger, William E. Weihl
November 1997 **ACM Transactions on Computer Systems (TOCS)**, Volume 15 Issue 4
Publisher: ACM Press
Full text available: pdf(259.35 KB) Additional Information: full citation, abstract, references, citings, index terms

This article describes the Digital Continuous Profiling Infrastructure, a sampling-based profiling system designed to run continuously on production systems. The system supports multiprocessors, works on unmodified executables, and collects profiles for entire systems, including user programs, shared libraries, and the operating system kernel. Samples are collected at a high rate (over 5200 samples/sec. per 333MHz processor), yet with low overhead (1–3% slowdown for most workloads). A ...

Keywords: performance understanding, performance-monitoring hardware, profiling, program analysis

7 The simulation of time sharing systems

Norman R. Nielsen
July 1967 **Communications of the ACM**, Volume 10 Issue 7

Publisher: ACM Press

Full text available: pdf(2.23 MB) Additional Information: full citation, abstract, references, citings, index terms

The development of new large scale time-sharing systems has raised a number of problems for computation center management. Not only is it necessary to develop an appropriate hardware configuration for these systems, but appropriate software adjustments must be made. Unfortunately, these systems often do not respond to changes in the manner that intuition would suggest, and there are few guides to assist in the analysis of performance characteristics. The development of a comprehensive simul ...

8 Partially preemptible hash joins

Hwee Hwa Pang, Michael J. Carey, Miron Livny

June 1993 **ACM SIGMOD Record , Proceedings of the 1993 ACM SIGMOD international conference on Management of data SIGMOD '93**, Volume 22 Issue 2

Publisher: ACM Press

Full text available: pdf(1.42 MB) Additional Information: full citation, abstract, references, citings, index terms

With the advent of real-time and goal-oriented database systems, priority scheduling is likely to be an important feature in future database management systems. A consequence of priority scheduling is that a transaction may lose its buffers to higher-priority transactions, and may be given additional memory when transactions leave the system. Due to their heavy reliance on main memory, hash joins are especially vulnerable to fluctuations in memory availability. Previous studies have propose ...

9 XML query processing I: Dynamic sample selection for approximate query processing

Brian Babcock, Surajit Chaudhuri, Gautam Das

June 2003 **Proceedings of the 2003 ACM SIGMOD international conference on Management of data**

Publisher: ACM Press

Full text available: pdf(260.80 KB) Additional Information: full citation, abstract, references, citings, index terms

In decision support applications, the ability to provide fast approximate answers to aggregation queries is desirable. One commonly-used technique for approximate query answering is sampling. For many aggregation queries, appropriately constructed biased (non-uniform) samples can provide more accurate approximations than a uniform sample. The optimal type of bias, however, varies from query to query. In this paper, we describe an approximate query processing technique that dynamically constructs ...

10 Using shared virtual memory for parallel join processing

Ambuj Shatdal, Jeffrey F. Naughton

June 1993 **ACM SIGMOD Record , Proceedings of the 1993 ACM SIGMOD international conference on Management of data SIGMOD '93**, Volume 22 Issue 2

Publisher: ACM Press

Full text available: pdf(1.08 MB) Additional Information: full citation, abstract, references, citings, index terms

In this paper, we show that shared virtual memory, in a shared-nothing multiprocessor, facilitates the design and implementation of parallel join processing algorithms that perform significantly better in the presence of skew than previously proposed parallel join processing algorithms. We propose two variants of an algorithm for parallel join processing using shared virtual memory, and perform a detailed simulation to investigate their performance. The algorithm is unique in that it employ ...

11 Architectures to make simple visualisations using simple systems

Alan Dix, Russell Beale, Andy Wood

May 2000 **Proceedings of the working conference on Advanced visual interfaces**

Publisher: ACM Press

Full text available: pdf(1.50 MB) Additional Information: full citation, abstract, references, citings, index terms

In previous work, the first author argued for simple lightweight visualisations. These are surprisingly complex to produce due to the need for infrastructure to read files, etc. onCue, a desktop 'agent', aids the rapid production of such visualisations and their integration with desktop and Internet applications. Two examples are used dancing histograms for 2D tables and pieTrees for hierarchical numeric data. A major focus is the importance of architecture, both that of onCue itself and th ...

Keywords: Internet—desktop integration, artificial intelligence, hierarchical data, interactive visualisation, software architecture

12 Online dynamic reordering

Vijayshankar Raman, Bhaskaran Raman, Joseph M. Hellerstein

December 2000 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 9 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(205.81 KB) Additional Information: full citation, abstract, citings, index terms

We present a pipelining, dynamically tunable *reorder* operator for providing user control during long running, data- intensive operations. Users can see partial results and accordingly direct the processing by specifying preferences for various data items; data of interest is prioritized for early processing. The reordering mechanism is efficient and non-blocking and can be used over arbitrary data streams from files and indexes, as well as continuous data feeds. We also investigate severa ...

Keywords: Informix, Interactive data processing, Online reordering, User control

13 The performance of multiversion concurrency control algorithms

Michael J. Carey, Waleed A. Muhanna

September 1986 **ACM Transactions on Computer Systems (TOCS)**, Volume 4 Issue 4

Publisher: ACM Press

Full text available: pdf(2.65 MB) Additional Information: full citation, abstract, references, citings, index terms

A number of multiversion concurrency control algorithms have been proposed in the past few years. These algorithms use previous versions of data items in order to improve the level of achievable concurrency. This paper describes a simulation study of the performance of several multiversion concurrency control algorithms, investigating the extent to which they provide increases in the level of concurrency and also the CPU, I/O, and storage costs resulting from the use of multiple versions. T ...

14 Computing curricula 2001

September 2001 **Journal on Educational Resources in Computing (JERIC)**

Publisher: ACM Press

Full text available: pdf(613.63 KB) html(2.78 KB) Additional Information: full citation, references, citings, index terms

15 File server scaling with network-attached secure disks

Garth A. Gibson, David F. Nagle, Khalil Amiri, Fay W. Chang, Eugene M. Feinberg, Howard Gobioff, Chen Lee, Berend Ozceri, Erik Riedel, David Rochberg, Jim Zelenka

June 1997 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1997 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '97**, Volume 25 Issue 1

Publisher: ACM Press

Full text available: pdf(1.77 MB) Additional Information: full citation, abstract, references, citings, index terms

By providing direct data transfer between storage and client, network-attached storage devices have the potential to improve scalability for existing distributed file systems (by removing the server as a bottleneck) and bandwidth for new parallel and distributed file systems (through network striping and more efficient data paths). Together, these advantages influence a large enough fraction of the storage market to make commodity network-attached storage feasible. Realizing the technology's ful ...

16 A study of I/O behavior of perfect benchmarks on a multiprocessor

A. L. Narasimha Reddy, Prithviraj Banerjee

May 1990 **ACM SIGARCH Computer Architecture News , Proceedings of the 17th annual international symposium on Computer Architecture ISCA '90**, Volume 18 Issue 3a

Publisher: ACM Press

Full text available: pdf(984.37 KB) Additional Information: full citation, abstract, references, citings, index terms

The I/O behavior of some scientific applications, a subset of Perfect benchmarks, executing on a multiprocessor is studied. The aim of this study is to explore the various patterns of I/O access of large scientific applications and to understand the impact of this observed behavior on the I/O subsystem architecture. I/O behavior of the program is characterized by the demands it imposes on the I/O subsystem. It is observed that implicit I/O or paging is not a major problem for the applicatio ...

17 RAID: high-performance, reliable secondary storage

Peter M. Chen, Edward K. Lee, Garth A. Gibson, Randy H. Katz, David A. Patterson

June 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 2

Publisher: ACM Press

Full text available: pdf(3.60 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Disk arrays were proposed in the 1980s as a way to use parallelism between multiple disks to improve aggregate I/O performance. Today they appear in the product lines of most major computer manufacturers. This article gives a comprehensive overview of disk arrays and provides a framework in which to organize current and future work. First, the article introduces disk technology and reviews the driving forces that have popularized disk arrays: performance and reliability. It discusses the tw ...

Keywords: RAID, disk array, parallel I/O, redundancy, storage, striping

18 Session 8: distributed systems: Transparent caching with strong consistency in dynamic content web sites

Cristiana Amza, Gokul Soundararajan, Emmanuel Cecchet

June 2005 **Proceedings of the 19th annual international conference on Supercomputing ICS '05**

Publisher: ACM Press

Full text available: pdf(370.07 KB) Additional Information: full citation, abstract, references

We consider a cluster architecture in which dynamic content is generated by a database back-end and a collection of Web and application server front-ends. We study the effect of transparent query caching on the performance of such a cluster. Transparency requires that cached entries be invalidated as a result of writes. We start with a coarse-grain table-level automatic invalidation cache. Based on observed workload characteristics, we enhance the cache with the necessary dependency tracking and ...

19 Query evaluation techniques for large databases

Goetz Graefe

June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

Publisher: ACM Press

Full text available: pdf(9.37 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

20 Scale and performance in a distributed file system

John H. Howard, Michael L. Kazar, Sherri G. Menees, David A. Nichols, M. Satyanarayanan, Robert N. Sidebotham, Michael J. West
February 1988 **ACM Transactions on Computer Systems (TOCS)**, Volume 6 Issue 1
Publisher: ACM Press

Full text available: pdf(2.38 MB) Additional Information: full citation, abstract, references, citings, index terms, review

The Andrew File System is a location-transparent distributed tile system that will eventually span more than 5000 workstations at Carnegie Mellon University. Large scale affects performance and complicates system operation. In this paper we present observations of a prototype implementation, motivate changes in the areas of cache validation, server process structure, name translation, and low-level storage representation, and quantitatively demonstrate Andrews ability to scale gracefully. W ...

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.
Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player